

IN THE CLAIMS:

Please amend claims 1, 16, 22, 28, 34, 39, 41, 44, 49, 54, and 69.

- 1 1. (Currently Amended) A method for a storage operating system implemented in a
2 storage system to optimize an amount of readahead data retrieved from
3 a data container of the storage system, the method comprising:
4 receiving a client read request for a particular read stream at the storage system;
5 ~~the client read request indicating client requested data for the storage operating system to~~
6 ~~retrieve from the data container~~;
7 locating a readset data structure for the particular read stream;
8 determining whether the storage operating system is permitted to retrieve reada-
9 head data from the data container in response to the received client read request;
10 if it is determined that the storage operating system is permitted to retrieve reada-
11 head data from the data container, performing the steps of:
12 (i) ~~locating one or more readset data structures associated with the client-~~
13 ~~requested data and establishing a read stream corresponding to each readset data structure~~
14 ~~of the one or more readset data structures~~;
15 (ii) (i) selecting an amount of readahead data to retrieve from the data container
16 based on a plurality of factors stored within a the readset data structure ~~of the one or more~~
17 ~~readset data structures~~ ; and
18 (iii) (ii) retrieving the selected amount of readahead data from the data container.
- 1 2. (Original) The method of claim 1, wherein the data container is a file, directory,
2 vdisk or lun.
- 1 3. (Original) The method of claim 1, wherein the storage operating system is deter-
2 mined to be permitted to retrieve readahead data from the data container when the client-
3 requested data extends the read stream past a predetermined next readahead value.

- 1 4. (Original) The method of claim 3, wherein the predetermined next readahead
2 value is stored in a readset data structure associated with the read stream.
- 1 5. (Original) The method of claim 3, wherein the predetermined next readahead
2 value is updated based on a percentage of the selected amount of readahead data.
- 1 6. (Previously Presented) The method of claim 1, wherein a read-access style asso-
2 ciated with the data container is one of the plurality of factors used to select the amount
3 of readahead data.
- 1 7. (Original) The method of claim 6, wherein the selected amount of readahead data
2 equals zero if the read-access style corresponds to a random read-access style.
- 1 8. (Previously Presented) The method of claim 1, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.
- 1 9. (Original) The method of claim 8, wherein the number of client read requests
2 processed in the read stream is stored as a count value in a readset data structure associ-
3 ated with the read stream.
- 1 10. (Previously Presented) The method of claim 1, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.
- 1 11. (Original) The method of claim 10, wherein the selected amount of readahead
2 data is set equal to a predetermined upper limit for large amounts of client-requested data.

1 12. (Original) The method of claim 1, wherein the selected amount of readahead data
2 is doubled if the number of client read requests processed in the read stream is greater
3 than a first threshold value.

1 13. (Original) The method of claim 1, wherein the client-requested data is identified
2 as read-once data when either (i) the number of client read requests processed in the read
3 stream is greater than a second threshold value or (ii) a set of metadata associated with
4 the read stream indicates that the client-requested data is read-once data.

1 14. (Original) The method of claim 1, wherein the selected amount of readahead data
2 is stored in one or more buffers enqueued on a flush queue, the flush queue being config-
3 ured to reuse buffers after a predetermined period of time.

1 15. (Original) The method of claim 14, wherein the predetermined period of time
2 equals two seconds.

1 16. (Currently Amended) An apparatus configured to implement a storage operating
2 system that optimizes an amount of readahead data retrieved from a data container of the
3 apparatus, the apparatus comprising:

4 means for receiving a client read request for a particular read stream, the client
5 read request indicating client requested data for the storage operating system to retrieve
6 from the data container;

7 means for locating a readset data structure for the particular read stream;

8 means for determining whether the storage operating system is permitted to re-
9 trieve readahead data from the data container in response to the received client read re-
10 quest;

11 means for locating one or more readset data structures associated with the client
12 requested data;

13 means for establishing a read stream corresponding to each readset data structure
14 of the one or more readset data structures;

15 means for selecting an amount of readahead data to retrieve from the data con-
16 tainer based on a plurality of factors stored within a the readset data structure of the one
17 or more readset data structures ; and

18 means for retrieving the selected amount of readahead data from the data con-
19 tainer.

1 17. (Original) The apparatus of claim 16, wherein the data container is a file, direc-
2 tory, vdisk or lun.

1 18. (Original) The apparatus of claim 16, wherein the storage operating system is de-
2 termined to be permitted to retrieve readahead data from the data container when the cli-
3 ent-requested data extends the read stream past a predetermined next readahead value.

1 19. (Original) The apparatus of claim 18, further comprising means for updating the
2 predetermined next readahead value based on a percentage of the selected amount of
3 readahead data.

1 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors
2 used to select the amount of readahead data includes at least one of:

- 3 (i) the amount of client-requested data,
- 4 (ii) a number of client read requests processed in the read stream, and
- 5 (iii) a read-access style associated with the data container.

1 21. (Original) The apparatus of claim 16, wherein the selected amount of readahead
2 data is doubled if the number of client read requests processed in the read stream is
3 greater than a first threshold value.

1 22. (Currently Amended) A storage system configured to optimize an amount of
2 readahead data retrieved from a data container of the storage system, the storage system
3 comprising:

4 a network adapter for receiving a client read request for a particular read stream;—
5 the client read request indicating client requested data to retrieve from the data container;
6 and

7 a memory configured to store instructions for implementing a storage operating
8 system that performs the steps of:

9 locating a readset data structure for the particular read stream; and
10 determining whether the storage operating system is permitted to retrieve reada-
11 head data from the data container in response to the received client read request;
12 and

13 if it is determined that the storage operating system is permitted to retrieve reada-
14 head data from the data container:

15 (i) locating one or more readset data structures associated with
16 the client requested data and establishing a read stream corre-
17 sponding to each readset data structure of the one or more readset
18 data structures;

19 (ii) selecting an amount of readahead data to retrieve from the data
20 container based on a plurality of factors stored within a readset
21 data structure of the one or more readset data structures; and

22 (iii) retrieving the selected amount of readahead data from the data container.

1 23. (Original) The storage system of claim 22, wherein the data container is a file,
2 directory, vdisk or lun.

1 24. (Original) The storage system of claim 22, wherein the storage operating system
2 is determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 25. (Original) The storage system of claim 24, wherein the predetermined next reada-
2 head value is updated based on a percentage of the selected amount of readahead data.

1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of
2 factors used to select the amount of readahead data includes at least one of:

3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 27. (Original) The storage system of claim 22, wherein the selected amount of reada-
2 head data is doubled if the number of client read requests processed in the read stream is
3 greater than a first threshold value.

1 28. (Currently Amended) A computer-readable media comprising instructions for
2 execution in a processor for the practice of a method for a storage operating system im-
3 plemented in a storage system to optimize an amount of readahead data retrieved from a
4 data container of the storage system, the method comprising:

5 receiving a client read request for a particular read stream at the storage system;
6 ~~the client read request indicating client requested data for the storage operating system to~~
7 ~~retrieve from the data container;~~

8 locating a readset data structure for the particular read stream;

9 determining whether the storage operating system is permitted to retrieve reada-
10 head data from the data container in response to the received client read request;

11 if it is determined that the storage operating system is permitted to retrieve reada-
12 head data from the data container, performing the steps of:

13 (i) ~~locating one or more readset data structures associated with the client-~~
14 ~~requested data and establishing a read stream corresponding to each read-~~
15 ~~set data structure of the one or more readset data structures;~~

16 (ii) (i) selecting an amount of readahead data to retrieve from the data container
17 based on a plurality of factors stored within a the readset data structure of the one or more
18 readset data structures ; and

19 (iii) (ii) retrieving the selected amount of readahead data from the data container.

1 29. (Original) The computer-readable media of claim 28, wherein the data container
2 is a file, directory, vdisk or lun.

1 30. (Previously Presented) The method of claim 1, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 31. (Previously Presented) The method of claim 30, wherein the read stream corre-
2 sponds to a read-once data transfer and data retrieved from the data container is stored in
3 the flush queue.

1 32. (Previously Presented) The method of claim 30, wherein the retrieved readahead
2 data is stored in the flush queue.

1 33. (Previously Presented) The method of claim 30, wherein one or more buffers ac-
2 cessed from the flush queue are re-enqueued on a normal queue.

1 34. (Currently Amended) A method for optimizing readahead data retrieved from
2 a data container of a storage system, the method comprising:
3 receiving a client read request for a particular read stream; client-requested data at
4 the storage system;
5 locating one or more a readset data structure structures for the particular read
6 stream; associated with the client-requested data;
7 establishing a read stream corresponding to each readset data structure of the one
8 or more readset data structures;
9 selecting an amount of readahead data based on the indicated amount of client-
10 requested data and a the corresponding readset data structure of the one or more readset
11 data structures ; and
12 retrieving the selected amount of readahead data from the data container.

1 35. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a multiple of a predetermined amount, and wherein the
3 multiple is associated with the amount of client-requested data.

1 36. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a multiple of the amount of client-requested data.

1 37. (Previously Presented) The method of claim 36, further comprising the step of
2 rounding the selected amount of readahead data to the size of a data block.

1 38. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a predetermined upper limit.

1 39. (Currently Amended) A method for optimizing readahead data retrieved from a
2 data container of a storage system, the method comprising:

3 receiving a client read request for a particular read stream at the storage system; ~~;~~
4 ~~the client read request indicating client requested data;~~

5 locating one or more a readset data structure structures ~~for the particular read~~
6 ~~stream; associated with the client requested data;~~

7 establishing a read stream corresponding to each readset data structure of the one
8 or more readset data structures;

9 selecting for the read stream an amount of readahead data based on a read-access
10 style associated with the data container and ~~a the~~ corresponding readset data structure of
11 the one or more readset data ; and

12 retrieving the selected amount of readahead data from the data container.

1 40. (Previously Presented) The method of claim 39, wherein the selected amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

1 41. (Currently Amended) A method for optimizing readahead data retrieval from
2 a data container of a storage system associated with a number of storage devices, the
3 method comprising:
4 receiving a client read request for a particular read stream at the storage system;
5 ~~the client read request indicating client requested data;~~
6 locating one or more a readset data structure structures for the particular read
7 stream; associated with the client requested data;
8 establishing a read stream corresponding to each readset data structure of the one
9 or more readset data structures;
10 selecting an amount of readahead data based on the number of storage devices
11 and a the corresponding readset data structure of the one or more readset data structures;
12 and
13 retrieving the selected amount of readahead data from the data container.

1 42. (Previously Presented) The method of claim 41, wherein the step of selecting an
2 amount of readahead data further comprises:
3 determining whether a flag is associated with the read stream, the flag indicating
4 that the storage system is associated with more than a predetermined number of storage
5 devices; and
6 in response to determining whether the flag is associated, selecting the amount of
7 readahead data.

1 43. (Previously Presented) The method of claim 41, wherein the storage devices
2 comprise one or more disks.

1 44. (Currently Amended) A method for optimizing readahead data retrieval in a stor-
2 age system, the method comprising:
3 receiving a client read request for a particular read stream at the storage system;
4 ~~for client requested data;~~

5 locating one or more a readset data structure structures for the particular read
6 stream; associated with the client requested data and establishing a read stream corre-
7 sponding to each readset data structure of the one or more readset data structures;
8 selecting an amount of readahead data based on a plurality of factors stored
9 within a the corresponding readset data structure of the one or more readset data struc-
10 tures ; and
11 retrieving the selected amount of readahead data from a data container.

1 45. (Previously Presented) The method of claim 44, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 46. (Previously Presented) The method of claim 45, wherein the read stream corre-
2 sponds to a read-once data transfer and data retrieved from the data container is stored in
3 the flush queue.

1 47. (Previously Presented) The method of claim 45, wherein the retrieved readahead
2 data is stored in the flush queue.

1 48. (Previously Presented) The method of claim 45, wherein one or more buffers ac-
2 cessed from the flush queue are re-enqueued on a normal queue.

1 49. (Currently Amended) A system for optimizing readahead data retrieval in a stor-
2 age system, the system comprising:
3 means for receiving a client read request for a particular read stream at the storage
4 system for client requested data;
5 means for locating one or more a readset data structure structures for the particu-
6 lar read stream associated with the client requested data and establishing a read stream
7 corresponding to each readset data structure of the one or more readset data structures;

8 means for selecting an amount of readahead data based on a plurality of factors
9 stored within a the corresponding readset data structure ~~of the one or more readset data~~
10 structures ; and

11 means for retrieving the selected amount of readahead data from ~~the~~ a data con-
12 tainer.

1 50. (Previously Presented) The system of claim 49, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 51. (Previously Presented) The system of claim 50, wherein the read stream corre-
2 sponds to a read-once data transfer and data retrieved from the data container is stored in
3 the flush queue.

1 52. (Previously Presented) The system of claim 50, wherein the retrieved readahead
2 data is stored in the flush queue.

1 53. (Previously Presented) The system of claim 50, wherein one or more buffers ac-
2 cessed from the flush queue are re-enqueued on a normal queue.

1 54. (Currently Amended) A method, comprising:
2 receiving a plurality of client read requests for a particular read stream at a storage
3 system; ~~the client read requests indicating client requested data sets for a storage operat-~~
4 ~~ing system to retrieve from one or more data containers;~~
5 locating one or more a readset data structure structures for the particular read
6 stream; ~~associated with the client requested data;~~
7 establishing a read stream corresponding to each readset data structure of the one
8 or more readset data structures;

9 selecting an amount of readahead data to retrieve from the one or more data con-
10 tainers based on a plurality of factors stored within a the readset data structure of the pla-
11 rality of readset data structures ;
12 retrieving the selected amount of readahead data from the data container;
13 processing one or more of the plurality of client read requests; and
14 adjusting, as client requests are processed, the plurality of factors stored within
15 the readset data structure associated with each read stream to optimize the amount of
16 readahead data that is cached for each read stream.

1 55. (Previously Presented) The method of claim 54, further comprising:
2 determining whether the storage operating system is permitted to retrieve reada-
3 head data from the one or more data containers in response to each received client read
4 request.

1 56. (Previously Presented) The method of claim 54, wherein the one or more data
2 containers are at least one of a file, a directory, a vdisk or a lun.

1 57. (Previously Presented) The method of claim 55, wherein the storage operating
2 system is determined to be permitted to retrieve readahead data from the one or more data
3 containers when the client-requested data extends the read stream past a predetermined
4 next readahead value.

1 58. (Previously Presented) The method of claim 57, wherein the predetermined next
2 readahead value is stored in a readset data structure associated with the read stream.

1 59. (Previously Presented) The method of claim 57, wherein the predetermined next
2 readahead value is updated based on a percentage of the selected amount of readahead
3 data.

- 1 60. (Previously Presented) The method of claim 54, wherein a read-access style as-
2 sociated with the one or more data containers is one of the plurality of factors used to se-
3 lect the amount of readahead data.

- 1 61. (Previously Presented) The method of claim 60, wherein the selected amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

- 1 62. (Previously Presented) The method of claim 54, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

- 1 63. (Previously Presented) The method of claim 62, wherein the number of client read
2 requests processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.

- 1 64. (Previously Presented) The method of claim 54, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.

- 1 65. (Previously Presented) The method of claim 64, wherein the selected amount of
2 readahead data is set equal to a predetermined upper limit for large amounts of client-
3 requested data.

- 1 66. (Previously Presented) The method of claim 54, wherein the selected amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

- 1 67. (Previously Presented) The method of claim 55, wherein the client-requested data
2 is identified as read-once data when either (i) the number of client read requests proc-

3 essed in the read stream is greater than a second threshold value or (ii) a set of metadata
4 associated with the read stream indicates that the client-requested data is read-once data.

1 68. (Previously Presented) The method of claim 54, wherein the selected amount of
2 readahead data is stored in one or more buffers enqueued on a flush queue, the flush
3 queue being configured to reuse buffers after a predetermined period of time.

1 69. (Currently Amended) A method for optimizing readahead data retrieval for a
2 storage system, the method comprising:

3 receiving a client read request for a particular read stream at the storage system;
4 for client-requested data stored in a file;

5 locating a plurality of readset data structure structures for the particular read
6 stream; associated with the client requested data;

7 establishing a read stream corresponding to each readset data structure of the plu-
8 rality of readset data structures;

9 selecting an amount of readahead data in response to a the readset data structure
10 of the plurality of readset data structures ; and

11 retrieving the selected amount of readahead data from the a data container.

1 70. (Previously Presented) The method of claim 69, further comprising:

2 allocating more readsets for the file in response to processing one or more client
3 “write” requests to the file.

1 71. (Previously Presented) A storage system, comprising:

2 a network adapter to receive a client read request at the storage system for client-
3 requested data stored in a file;

4 a plurality of readset data structures associated with the client-requested data;

5 an operating system to establish a read stream corresponding to each readset data
6 structure of the plurality of readset data structures;

7 a process to select an amount of readahead data in response to a readset data
8 structure of the plurality of readset data structures; and
9 a storage adapter to retrieve the selected amount of readahead data from the data
10 container.

1 72. (Previously Presented) The storage system of claim 71, further comprising:
2 the operating system to allocate more readsets for the file in response to process-
3 ing one or more client “write” requests to the file.